

Report Documentation Page		Form Approved OMB No. 0704-0188
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.		
1. REPORT DATE 13 JUN 2013	2. REPORT TYPE Final	3. DATES COVERED 30-06-2010 to 29-06-2013
4. TITLE AND SUBTITLE Extending Semantic and Episodic Memory to Support Robust Decision Making		5a. CONTRACT NUMBER FA23861014127
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S) John Laird		5d. PROJECT NUMBER
		5e. TASK NUMBER
		5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The Regents of the University of Michigan,3003 S. State Street,Ann Arbor,MI,48109-1287		8. PERFORMING ORGANIZATION REPORT NUMBER N/A
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AOARD, UNIT 45002, APO, AP, 96338-5002		10. SPONSOR/MONITOR'S ACRONYM(S) AOARD
		11. SPONSOR/MONITOR'S REPORT NUMBER(S) AOARD-104127
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited		
13. SUPPLEMENTARY NOTES		

14. ABSTRACT

Our research focused on developing and evaluating general, effective, and efficient algorithms for learning of long-term knowledge in autonomous agents, as well as developing cognitive capabilities that exploit that learning. Our work has covered episodic memory, semantic memory, and procedural memory, integrated within a general cognitive architecture (Soar). For episodic memory, our research has led to significant improvements in the efficiency of storage (memory) and retrieval (time) through the exploitation of temporal contiguity, structural regularity, high cue structural selectivity, high temporal selectivity, low cue feature co-occurrence, resulting in no significant slowdown with experience: runs for days of real time (tens of millions of episodes), faster than real time [Laird, Derbinsky, & Voigt 2011; Derbinsky, Li & Laird 2012a]. We evaluated our approach on multiple tasks (including mobile robotics, games, planning problems, linguistics) and for multiple cognitive capabilities [Laird, Derbinsky, & Voigt 2011; Laird & Derbinsky 2011c; Derbinsky, Li & Laird 2012b]. For semantic memory retrieval, we studied and evaluated multiple functions for biasing retrieval, and developed efficient approximate algorithm that maintains high (>90%) validity for base-level activation (the best performing in terms of quality). Our algorithm is 30-100 times faster than prior retrieval algorithms and has sub linear slowdown as memory size increases [Derbinsky, Laird, & Smith 2010; Laird & Derbinsky 2011a; Derbinsky & Laird 2011b]. We also developed algorithms to use semantic memory to support reconstruction of forgotten working memory [Derbinsky & Laird 2011c; Derbinsky & Laird 2012a]. This research was evaluated on multiple domains including word sense disambiguation and mobile robotics. [Derbinsky, Laird, & Smith 2010; Laird, Derbinsky, & Voigt 2011; Derbinsky & Laird 2011b]. We also developed algorithms for retaining/forgetting of knowledge in working memory and procedural memory to improve the scalability of episodic memory and procedural memory [Laird & Derbinsky 2011c; Derbinsky & Laird 2012b; Derbinsky & Laird 2013]. Our final research was on learning value functions for reinforcement learning [Laird, Derbinsky, Tinkerhess 2012; Bloch & Laird 2013].

15. SUBJECT TERMS

Artificial Intelligence, Autonomous Agents and Multi-Agent Systems, Machine Learning

16. SECURITY CLASSIFICATION OF:

a. REPORT

unclassified

b. ABSTRACT

unclassified

c. THIS PAGE

unclassified17. LIMITATION OF
ABSTRACT**Same as
Report (SAR)**18. NUMBER
OF PAGES**2**19a. NAME OF
RESPONSIBLE PERSON

"Extending Semantic and Episodic Memory to Support Robust Decision Making"

June 13, 2013

Name of Principal Investigators: John E. Laird

- e-mail address : laird@umich.edu
- Institution : University of Michigan
- Mailing Address : 2260 Hayward Street, Ann Arbor, MI 48109-2121
- Phone : 734 647-1761
- Fax : 734 763-1260

Period of Performance: 7/1/2010-6/30/2013

Abstract: Our research focused on developing and evaluating general, effective, and efficient algorithms for learning of long-term knowledge in autonomous agents, as well as developing cognitive capabilities that exploit that learning. Our work has covered episodic memory, semantic memory, and procedural memory, integrated within a general cognitive architecture (Soar). For episodic memory, our research has led to significant improvements in the efficiency of storage (memory) and retrieval (time) through the exploitation of temporal contiguity, structural regularity, high cue structural selectivity, high temporal selectivity, low cue feature co-occurrence, resulting in no significant slowdown with experience: runs for days of real time (tens of millions of episodes), faster than real time [Laird, Derbinsky, & Voigt 2011; Derbinsky, Li & Laird 2012a]. We evaluated our approach on multiple tasks (including mobile robotics, games, planning problems, linguistics) and for multiple cognitive capabilities [Laird, Derbinsky, & Voigt 2011; Laird & Derbinsky 2011c; Derbinsky, Li & Laird 2012b].

For semantic memory retrieval, we studied and evaluated multiple functions for biasing retrieval, and developed efficient approximate algorithm that maintains high (>90%) validity for base-level activation (the best performing in terms of quality). Our algorithm is 30-100 times faster than prior retrieval algorithms and has sub linear slowdown as memory size increases [Derbinsky, Laird, & Smith 2010; Laird & Derbinsky 2011a; Derbinsky & Laird 2011b]. We also developed algorithms to use semantic memory to support reconstruction of forgotten working memory [Derbinsky & Laird 2011c; Derbinsky & Laird 2012a]. This research was evaluated on multiple domains including word sense disambiguation and mobile robotics. [Derbinsky, Laird, & Smith 2010; Laird, Derbinsky, & Voigt 2011; Derbinsky & Laird 2011b].

We also developed algorithms for retaining/forgetting of knowledge in working memory and procedural memory to improve the scalability of episodic memory and procedural memory [Laird & Derbinsky 2011c; Derbinsky & Laird 2012b; Derbinsky & Laird 2013].

Our final research was on learning value functions for reinforcement learning [Laird, Derbinsky, Tinkerhess 2012; Bloch & Laird 2013].

List of Publications and Significant Collaborations that resulted from your AOARD supported project

a) papers published in peer-reviewed journals,

1. Derbinsky, N., and Laird, J. E. (2013) Effective and efficient forgetting of learned knowledge in Soar's working and procedural memories. *Cognitive Systems Research*, 24, 104-113.
2. Laird, J. E., Derbinsky, N. and Tinkerhess, M. (2012). Online Determination of Value-Function Structure and Action-value Estimates for Reinforcement Learning in a Cognitive Architecture, *Advances in Cognitive Systems*, Volume 2, 221-238, Palo Alto, California.

b) papers published in peer-reviewed conference proceedings,

3. Derbinsky, N., Li, J., and Laird, J. E. (2012a) Algorithms for Scaling in a General Episodic Memory (Extended Abstract). *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. Valencia, Spain.

4. Derbinsky, N., and Laird, J. E. (2012a) Computationally Efficient Forgetting via Base-Level Activation. Proceedings of the 11th International Conference on Cognitive Modeling (ICCM). 109-110. Berlin, Germany. *Best Poster*.
5. Derbinsky, N., and Laird, J. E. (2012b) Competence-Preserving Retention of Learned Knowledge in Soar's Working and Procedural Memories. Proc. of the 11th International Conference on Cognitive Modeling (ICCM). 205-210. Berlin, Germany.
6. Derbinsky, N., Li, J., and Laird, J. E. (2012b) A Multi-Domain Evaluation of Scaling in a General Episodic Memory. Proceedings of the 26th AAAI Conference on Artificial Intelligence. 193-199. Toronto, Canada.
7. Li, J., Derbinsky, N., and Laird, J. E. (2012) Functional Interactions Between Encoding and Recognition of Semantic Knowledge. Proceedings of the 26th AAAI Conference on Artificial Intelligence. 228-234 Toronto, Canada.
8. Laird, J. E., Derbinsky, N., and Voigt, J. (2011) Performance Evaluation of Declarative Memory Systems in Soar (2011). Proceedings of the 20th Behavior Representation in Modeling & Simulation Conference (BRIMS), 33-40. Sundance, UT.
9. Derbinsky, N., and Laird, J. E. (2011a) A Preliminary Functional Analysis of Memory in the Word Sense Disambiguation Task. Proceedings of the 2nd Symposium on Human Memory for Artificial Agents, AISB, 25-29. York, England.
10. Derbinsky, N., and Laird, J. E. (2011b) A Functional Analysis of Historical Memory Retrieval Bias in the Word Sense Disambiguation Task. Proceedings of the 25th National Conference on Artificial Intelligence (AAAI). 663-668. San Francisco, CA.
11. Derbinsky, N., and Laird, J. E. (2011c) Effective and Efficient Management of Soar's Working Memory via Base-Level Activation. Papers from the 2011 AAAI Fall Symposium Series: Advances in Cognitive Systems (ACS), 82-89. Arlington, VA.
12. Derbinsky, N., Laird, J. E., and Smith, B. (2010) Towards Efficiently Supporting Large Symbolic Declarative Memories. Proc. of the 10th International Conference on Cognitive Modeling (ICCM), 49-54. Philadelphia, PA.

c) papers published in non-peer-reviewed journals and conference proceedings: none

d) conference presentations without papers: none

e) manuscripts submitted but not yet published:

Bloch, M., and Laird, J. E. (submitted) Online Value Function Improvement, The First Multidisciplinary Conference on Reinforcement Learning and Decision Making RLDM 2013.

f) interactions with industry or with Air Force Research Laboratory scientists or significant collaborations that resulted from this work.

- We have been in many discussions with AFRL researchers from the PALM research group at Wright Patterson Air Force Base (led by Kevin Gluck) about our research. Our research on semantic memory was inspired in part by research they had done on scaling of long-term declarative memory.
- This work also has been used in other research projects, including part of DARPA BOLT section E, and my two ONR projects managed by Paul Bello.
- Our research is also used in industry, specifically Soar Technology (episodic and semantic memory) and L3 Communications (reinforcement learning).

Attachments: Publications a), b) and c) listed above if possible.

DD882: As a separate document, please complete and sign the inventions disclosure form.